AEDC-TR-86-33
C.2

Second-Generation Discrete Fourier Transform Signal Processor for Laser Velocimetry

Thomas Chris Layne
Calspan Corporation
and
Dr. Bruce William Bomar
University of Tennessee Space Institute

December 1986

Final Report for Period October 1, 1984 – August 15, 1986

ARNOLD ENGINEERING DEVELOPMENT CENTER
ARNOLD AIR FORCE STATION, TENNESSEE
AIR FORCE SYSTEMS COMMAND
UNITED STATES AIR FORCE

NOTICES

When U. S. Government drawings, specifications, or other data are used for any purpose other than a definitely related Government procurement operation, the Government thereby incurs no responsibility nor any obligation whatsoever, and the fact that the Government may have formulated, furnished, or in any way supplied the said drawings, specifications, or other data, is not to be regarded by implication or otherwise, or in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

Qualified users may obtain copies of this report from the Defense Technical Information Center.

References to named commercial products in this report are not to be considered in any sense as an endorsement of the product by the United States Air Force or the Government.

This report has been reviewed by the Office of Public Affairs (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nations.

APPROVAL STATEMENT

This report has been reviewed and approved.

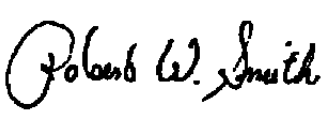

ROBERT W. SMITH
Directorate of Technology
Deputy for Operations

Approved for publication:

FOR THE COMMANDER

LOWELL C. KEEL, Colonel, USAF
Director of Technology
Deputy for Operations

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION UNCLASSIFIED	1b. RESTRICTIVE MARKINGS
2a SECURITY CLASSIFICATION AUTHORITY	3 DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution is unlimited.
2b DECLASSIFICATION/DOWNGRADING SCHEDULE	
4 PERFORMING ORGANIZATION REPORT NUMBER(S) AEDC-TR-86-33	5. MONITORING ORGANIZATION REPORT NUMBER(S)
6a NAME OF PERFORMING ORGANIZATION Arnold Engineering Development Center 6b. OFFICE SYMBOL (If applicable) DOT	7a NAME OF MONITORING ORGANIZATION
6c ADDRESS (City, State and ZIP Code) Air Force Systems Command Arnold Air Force Station, TN 37389-5000	7b. ADDRESS (City, State and ZIP Code)
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Arnold Engineering Development Center 8b. OFFICE SYMBOL (If applicable) DO	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER
8c ADDRESS (City, State and ZIP Code) Air Force Systems Command Arnold Air Force Station, TN 37389-5000	10 SOURCE OF FUNDING NOS. PROGRAM ELEMENT NO.: 65807F PROJECT NO TASK NO WORK UNIT NO
11. TITLE (Include Security Classification) SEE REVERSE OF THIS PAGE	
12. PERSONAL AUTHOR(S) Layne, Thomas Chris, Calspan Corporation, AEDC Division and Bomar, Bruce William, (cont)	
13a. TYPE OF REPORT Final 13b TIME COVERED FROM 10/1/84 TO 8/15/86	14 DATE OF REPORT (Yr., Mo., Day) December 1986 15. PAGE COUNT 40
16. SUPPLEMENTARY NOTATION Available in Defense Technical Information Center (DTIC).	

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB. GR.	
14	02		digital signal processing; windowing; laser velocimeter signal; signal conditioning; discrete Fourier transform; electro-optics
12	01		

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

A high-speed, single-board digital signal processor has been developed at the AEDC for frequency domain analysis of laser velocimeter signals. The analysis is facilitated by the discrete Fourier transform (DFT), which is used to determine a signal's spectral content. The processor consists of a microprogrammable microprocessor controlling circuitry specially designed for fast, efficient computation of the DFT. The DFT computation circuitry, which features a dual multiplier/accumulator architecture, allows computation of one DFT coefficient per 11.3 μsec for 113 data samples. The microprocessor searches for a dominant maximum in the frequency spectrum corresponding to the frequency of the input signal burst. Results from a laboratory evaluation of the processor are also included.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT ☒ DTIC USERS ☐	21 ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED
22a NAME OF RESPONSIBLE INDIVIDUAL W. O. Cole	22b TELEPHONE NUMBER (Include Area Code) (615)454-7813 22c OFFICE SYMBOL DOS

DD FORM 1473, 83 APR EDITION OF 1 JAN 73 IS OBSOLETE.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

11. TITLE

Second-Generation Discrete Fourier Transform Signal Processor for Laser Velocimetry

12. PERSONAL AUTHORS (Concluded)

University of Tennessee Space Institute

PREFACE

The work reported herein was performed at the Arnold Engineering Development Center (AEDC), Air Force Systems Command (AFSC). The AEDC Project Manager was Mr. Robert W. Smith. The results were obtained by Calspan Corporation, AEDC Division, operating contractor for the aerospace flight dynamics testing facilities at the AEDC, AFSC, Arnold Air Force Station, Tennessee, under Project Number D283 (V32N-B3). The research was conducted from October 1, 1984 through August 15, 1986, and the manuscript was submitted for publication on August 28, 1986.

CONTENTS

ILLUSTRATIONS

TABLE

1.0 INTRODUCTION

In 1983, Kalb and Crosswy reported the development of a prototype laser velocimeter (LV) signal processor that employs spectral analysis to determine input signal frequency (Ref. 1). Prior to the advent of this processor, signal processors using powerful spectral analysis techniques were too slow for most laser velocimetry applications. Kalb's processor uses a high-speed waveform digitizer to capture signal bursts and specially designed computing circuits to perform spectral analysis via the discrete Fourier transform (DFT). Therefore, the prototype processor was referred to as the DFT processor.

Successful application of the prototype DFT processor in Aerodynamic Wind Tunnel (1T) (Ref. 2) and Supersonic Wind Tunnel (A) (Ref. 3) gave rise to an effort to develop a reliable, maintainable version of the DFT processor. This report describes, in detail, the production model DFT processor that resulted. Theory of operation of the DFT processor's hardware and signal processing algorithms are presented. The results of a laboratory evaluation of the processor are also included.

1.1 FUNDAMENTAL CONSIDERATIONS

The laser velocimeter was developed to provide a nonintrusive technique for fluid velocity measurement. In the popular dual-beam type LV optical system shown in Fig. 1, two mutually coherent and polarization-aligned laser beams are crossed to form interference fringes having a known spacing. This crossover region defines a region in space called the probe volume. When a particle entrained in the flow passes through the probe volume, it is alternately illuminated by the maximum and minimum intensity of the interference fringes. Any light that the particle reflects out of the probe volume is focused onto the photodetector by the receiving lens assembly. The output of the photodetector is an amplitude-fluctuating signal whose rate of amplitude fluctuation (frequency) is proportional to the velocity of the particle traveling through the probe volume.

Presently, the counter-type processor is the most widely used means of extracting particle velocity information from LV signals. The counter-type processor employs a level-crossing detector and a high-speed timer to determine signal frequency. The level-crossing detector converts the analog LV signal into a digital pulse train, and the high-speed timer measures the time required for some preset number of level-crossings to occur. Because the counter processor's measurements are based directly on the time domain representation of the signal burst, an input signal that is relatively noise-free and well-defined is required for an accurate counter-processor measurement.

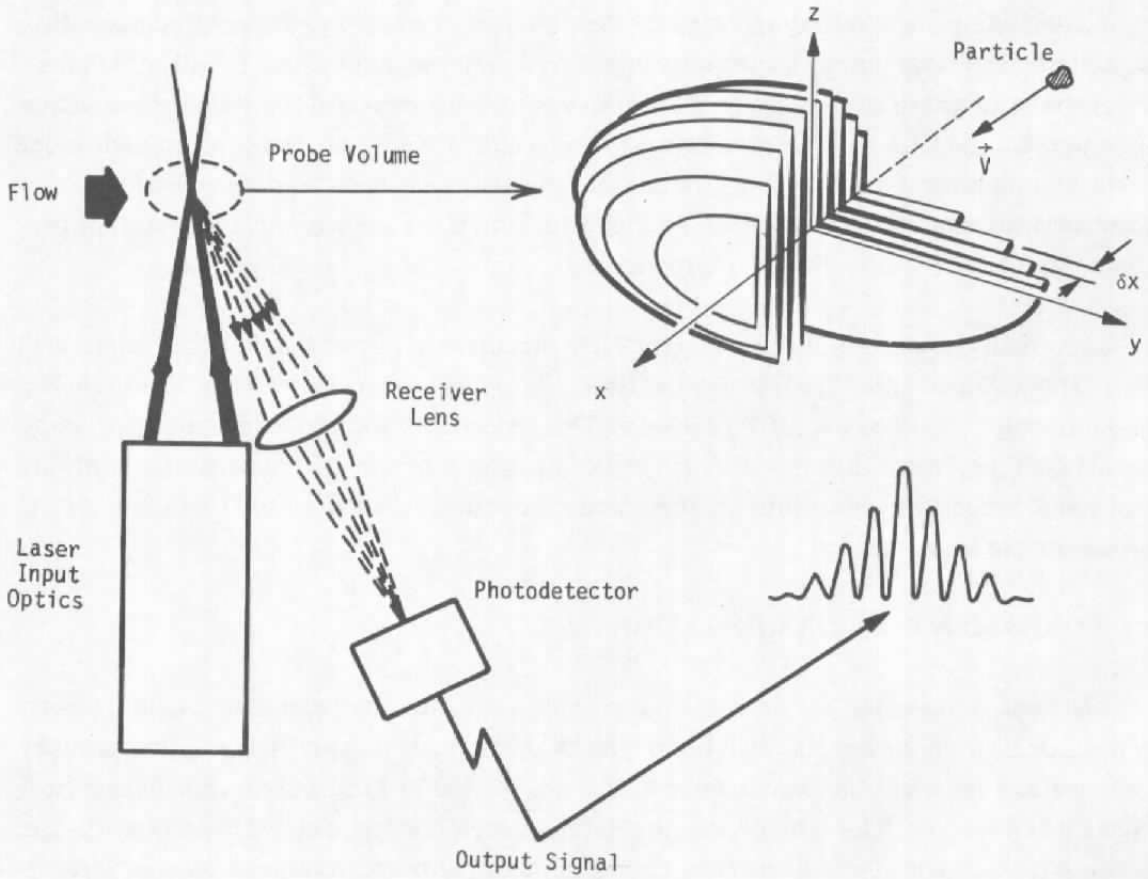

Figure 1. Dual-beam laser velocimeter.

Since fluid velocity information is inferred from particle velocity measurements, the accuracy of LV data is limited by the particles' ability to track the fluid velocity. Measurements in high fluid velocity gradients require seed particles less than 1 μm in diameter to satisfy particle dynamics constraints. Such small particles generate ill-defined signals often characterized by low signal-to-noise ratios. Unfortunately, it is difficult to process signals of this type using conventional counter-type signal processors.

Even though spectral analysis of ill-defined LV signals has long been considered superior to time-domain analysis, early frequency domain processors did not find widespread application because of their relatively long data acquisition time. The development of certain high-speed very-large-scale-integration (VLSI) integrated circuits enabled Kalb (Ref. 1) to design a fast

frequency domain LV signal processor based on waveform digitization and spectral analysis via the discrete Fourier transform. Kalb's prototype DFT processor was incorporated into the data acquisition system shown in Fig. 2. The transient digitizer samples LV signals and transfers signal samples to the DFT processor for analysis. The DFT processor uses a microprogrammable microprocessor to control circuitry specially designed for fast fixed-point-arithmetic calculation of the DFT. Spectral data corresponding to particle velocity measurements are stored in the host computer system. Kalb's DFT processor could perform 300 spectral analyses per second, an acceptable data rate for many applications. Tests in Tunnel 1T and Tunnel A demonstrated the ability of the prototype DFT processor to extract meaningful information from poor-quality signals that were rejected by a counter-type processor.

1.2 DESIGN PHILOSOPHY

The goal of this project was to transform the prototype DFT processor into a tool for routine use in the acquisition of LV data. Therefore, certain design criteria were established. Since any reduction in data acquisition time reduces tunnel time required for a given test, minimizing waveform processing time is essential. Because processor downtime is equivalent to tunnel downtime during a test, processor reliability and maintainability are also important. Given this, speed, reliability, and maintainability were emphasized during the design of the production model DFT.

The factor most limiting the throughput of the prototype processor is the inordinate amount of time required to transfer digitized waveforms from the Tektronix 7612D digitizer to the processor. The prototype processor is capable of processing 960 waveforms per second; unfortunately, the digitizer's maximum waveform transfer rate is only 300 waveforms per second. Therefore, the prototype processor is constrained to 300 measurements per second. A digitizer capable of transferring 18,000 waveforms per second was designed for use with the production model DFT processor. Therefore, the new DFT processor is limited only by its waveform processing speed.

When processing LV signals, the DFT processor must first compute a frequency spectrum via the discrete Fourier transform, then it must locate a dominant spectral peak of the transform that corresponds to a particle velocity. The most important factor in determining the time required to compute a spectrum is the speed of the microprocessor-controlled DFT computation hardware. The amount of time required for peak location is primarily determined by the efficiency of the signal processing algorithm. Thus, maximum processing speed is obtained by tailoring the architecture of the hardware to rapidly computing frequency spectra and by programming a computationally efficient peak location algorithm into the on-board microprocessor.

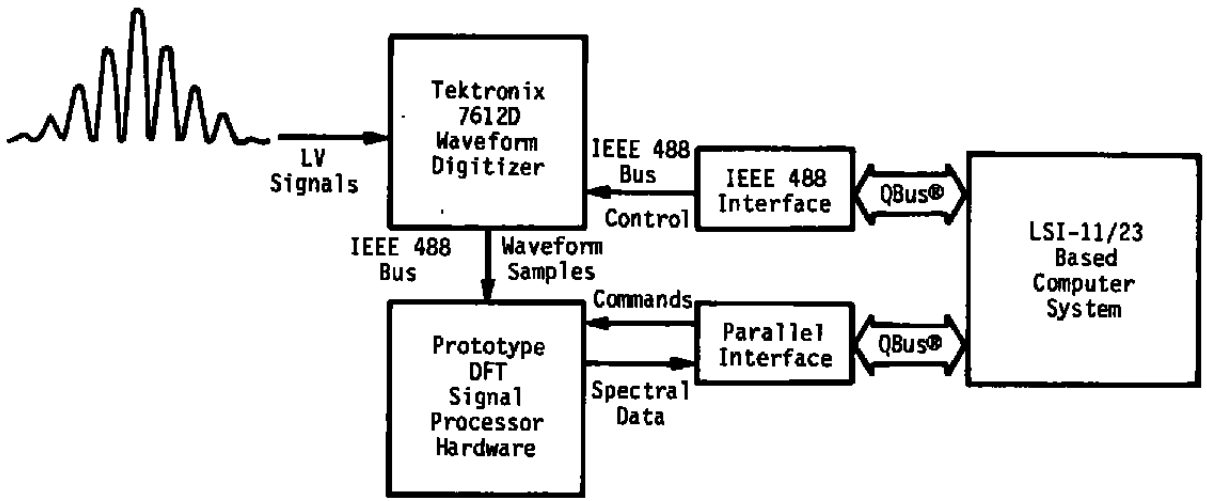

Figure 2. Prototype DFT processor LV data acquisition system.

The single most important factor affecting the reliability of the prototype DFT processor is the number of wired interconnections it contains. Jumper wires interconnect the various circuit boards that comprise the prototype processor. Cabling is also used to connect the processor, host computer, and digitizer. The production DFT processor was designed to eliminate as many of these wired interconnects as possible. Jumper wires that were used to connect integrated circuits were eliminated by building the new processor on a single, four-layer printed circuit board. Also, cabling from the processor to the host computer interface was eliminated by including the host interface circuitry on the processor board. The production model DFT processor resides on a single printed circuit board that plugs directly into and operates as an integral part of the host computer system.

A complex digital system consisting of 61 integrated circuits and over 2,000 separate interconnections resides on the processor's four-layer printed circuit board. Generally speaking, maintaining and troubleshooting such a complicated system is a formidable task. However, the difficulty for this task was greatly reduced by including several self-test diagnostic routines in the on-board microprocessor's program memory. The processor self-test is performed at power-up and each time the processor is reset. Should the processor fail to successfully complete any of the diagnostic routines, data acquisition is prohibited, and a light emitting diode (LED) is activated to indicate a processor fault. At this point, a functional processor board can be quickly substituted for the faulty one, allowing data acquisition to resume.

A diagnostic board was also developed to facilitate the repair of faulty DFT processor boards. The diagnostic board plugs onto a special port on the DFT board and monitors the operation of the DFT processor's on-board microprocessor. The current instruction being executed as well as the address of the next instruction are displayed in LED readouts on the diagnostic board. The on-board microprocessor's master clock may be single-stepped via the diagnostic board, allowing static faults to be easily located. The diagnostic board also contains an array of erasable-programmable read-only memory (EPROM), which may be configured to serve as the on-board microprocessor's program memory.

2.0 PROCESSING ALGORITHM

In digital signal processing, signals are represented by sequences of numbers. For laser velocimetry the digital signal of interest is obtained by using an analog-to-digital converter (A/D) to sample the output, p(t), of a photodetector. The resulting digital signal is a sequence of numbers p(n) for $n = 0, 1, ..., N - 1$. The spacing, Δt, of the samples in this sequence is the reciprocal of the A/D sample rate, f_s. The time interval $T = N\Delta t$ over which the sequence is defined is referred to as the signal window.

As shown in Fig. 3, the DFT processor processes the digital signal p(n) in three steps. The first step is to remove the pedestal and d-c components from the signal with a high-pass trend removal filter. The second step is to gradually taper the leading and trailing edges of the signal with a window function to prevent spectral distortion attributable to signal truncation in the time-domain window (Ref. 1). The third step is to calculate enough of the signal spectrum to permit an accurate location of the spectral peak, and hence, an accurate determination of the signal frequency. In this step the presence of a dominant spectral peak is required to validate the spectrum and associated measurement. Measurements that can not be validated are rejected.

2.1 TREND REMOVAL

To remove the pedestal and d-c frequency components from the signal, a high-pass finite-impulse-response (FIR) digital filter is employed. The filter output, q(n), is obtained from p(n) using the filter equation.

$$q(n - 8) = 16\, p(n) - \sum_{\ell = -8}^{7} p(n + \ell),\ n = 8, 9, \ldots, N - 8 \qquad (1)$$

Thus, each sample is filtered by multiplying it by 16 and then removing from it the sum, or average, of the 16 samples around the filtered sample.

For every sixteenth sample of p(n), the filter of Eq. (1) is equivalent to the filter employed in the prototype DFT processor (Ref. 1). However, between these samples the prototype DFT processor would use linear interpolation to estimate the mean that should be removed, rather than actually calculating the mean as is done in the new DFT processor.

The filter of Eq. (1) is efficiently implemented in the new DFT processor by noting that q(n) can be written as

$$m(8) = \sum_{\ell = 0}^{15} p(\ell) \qquad (2)$$

$$\left.\begin{aligned} q(n - 8) &= 16\, p(n) - m(n) \\ m(n + 1) &= m(n) + p(n + 8) - p(n - 8) \end{aligned}\right\} \quad n = 8, 9, \ldots, N - 8 \qquad (3)$$

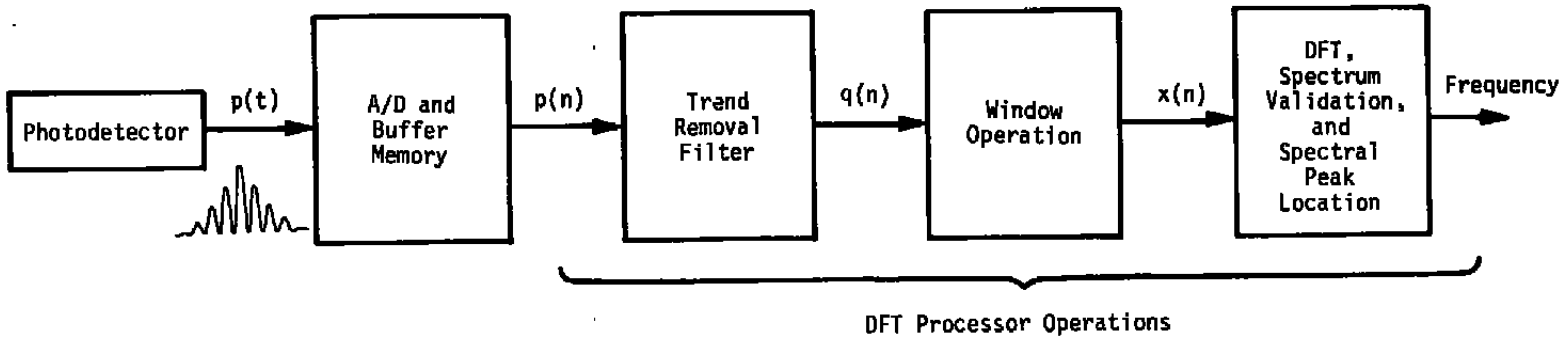

Figure 3. Steps in processing the photodetector signal.

Therefore, after initialization by Eq. (2) the filter can be implemented by Eq. (3), which requires only a shift, one addition, and two subtractions per sample filtered. After filtering, only 113 of the original 128 signal samples remain. Fifteen samples are lost in the filtering process in order to avoid filter transients.

2.2 WINDOWING

The window function chosen for use in the new DFT processor is the Hamming window. This window provides a main-lobe 3-db bandwidth, which is 9 percent less than that of the Hanning window used in the prototype DFT processor. This reduction in main-lobe bandwidth provides a somewhat higher spectral resolution. Also, the side-lobe level for the Hamming window is 10 db lower than for the Hanning window. The lower side-lobe level produces less spectral distortion if multiple frequency components are present in the signal.

The equation for the Hamming window samples is

$$w(n) = 0.54 - 0.46 \cos\left(\frac{2\pi n}{N - 16}\right), \; n = 0, 1, \ldots, N - 16 \tag{4}$$

The windowed data samples are then given by

$$x(n) = w(n)q(n), \; n = 0, 1, \ldots, L - 1 \tag{5}$$

where $L = N - 15$.

2.3 SIGNAL SPECTRUM

The continuous Fourier transform of the trend-removed and windowed signal x(n) is given by (Ref. 4)

$$X(e^{j\omega}) = \sum_{n=0}^{L-1} x(n)e^{-j\omega n} \tag{6}$$

This spectrum is a periodic function of digital radian frequency ω. The fundamental portion of the spectrum between $\omega = 0$ and $\omega = 2\pi$ repeats periodically with period 2π. Therefore, the inverse Fourier transform can be defined in terms of only the fundamental spectrum giving

$$x(n) = \frac{1}{2\pi} \int_0^{2\pi} X(e^{j\omega}) \, e^{j\omega n} d\omega \tag{7}$$

For digital signals, only frequencies from d-c through half the sample rate can be unambiguously determined. Thus, the Nyquist frequency $f_N = f_s/2$ corresponds to the highest processable signal frequency. In terms of digital radian frequency, $\omega = 0$ corresponds to d-c and $\omega = \pi$ corresponds to the Nyquist frequency. Therefore, a frequency, f, in the A/D input signal is related to ω in the spectrum of Eq. (6) through

$$\omega = \frac{f}{(f_s/2)} \pi \tag{8}$$

Values of ω between π and 2π correspond to negative values of f and are unimportant for real-valued input signals.

If we evaluate the fundamental spectrum of Eq. (6) on the L uniformly spaced digital radian frequencies,

$$\omega_k = \frac{2\pi}{L} k, \; k = 0, 1, \ldots, L - 1 \tag{9}$$

we obtain

$$X(k) \triangleq X(e^{j\omega_k}) = \sum_{n=0}^{L-1} x(n) \, e^{-j\frac{2\pi}{L}kn}, \; k = 0, 1, \ldots, L - 1 \tag{10}$$

Equation (10) is referred to as the discrete Fourier transform (DFT) of x(n). The DFT has the property (Ref. 4) that the time-domain samples x(n) can be uniquely recovered from the frequency-domain coefficients X(k) using the inverse DFT equation.

$$x(n) = \frac{1}{L} \sum_{k=0}^{L-1} X(k) \, e^{j\frac{2\pi}{L}kn}, \; n = 0, 1, \ldots, L - 1 \tag{11}$$

In processing laser velocimeter signals, we are only interested in the positive-frequency spectrum, which corresponds to $\omega = 0$ to $\omega = \pi$. Furthermore, we are not interested in inverting the spectrum to recover x(n). Therefore, we are not constrained to evaluate Eq. (6) on the particular grid of frequencies specified by Eq. (9). Instead, we will use the grid of frequencies

$$\omega_k = \frac{2\pi}{M} k, \; k = 0, 1, \ldots, \frac{M}{2} \tag{12}$$

giving

$$X(k) = \sum_{n=0}^{L-1} x(n)\, e^{-j\frac{2\pi}{M}kn},\ k = 0, 1, \ldots, \frac{M}{2} \qquad (13)$$

By choosing $M >> L$ we can obtain a spectrum with much finer resolution than would be possible by simply computing the DFT. Since the positive-frequency DFT coefficients are obtained from Eq. (13) when $M = L$, we will still refer to the coefficients calculated by Eq. (13) as "DFT" coefficients, although this is only strictly true when $M = L$. When $M >> L$, Eq. (13) really gives the coefficients of a zero-padded DFT.

2.4 LOCATING AND VALIDATING THE SPECTRAL PEAK

The peak-location algorithm presently microprogrammed on the new DFT processor was developed under the following assumptions:

1. It is desirable to use all A/D samples in computing each DFT coefficient.

2. The frequency tracking algorithm employed in the prototype processor computes only those DFT coefficients that lie in its narrow tracking band. This type of algorithm should be avoided because of possible problems in turbulent flow.

3. The number of DFT coefficients computed should be minimized to minimize processing time.

4. Each LV signal will be digitized as $N = 128$ samples spaced $\Delta t = 10$ nsec apart. This gives a Nyquist frequency of 50 MHz and a value of $L = 113$.

5. For hardware considerations the value of M used in Eq. (13) will be $M = 4{,}096$.

To avoid frequency tracking in the new DFT processor, DFT coefficients spanning the entire frequency range of interest are computed for each signal. However, to minimize the number of DFT coefficients that must be determined, these coefficients are spaced as far apart as possible without losing a spectral peak between two coefficients. This is accomplished by making the coefficient spacing approximately equal to the 3-db bandwidth of the main lobe in the Fourier transform of the Hamming window. This 3-db bandwidth is given by (Ref. 5)

$$B_{HW} = \frac{1.26}{L\Delta t} \qquad (14)$$

and so we require

$$\text{Coefficient Spacing} \approx \frac{1.26}{L\Delta t} \tag{15}$$

For the typical values of N = 128 and Δt = 10 nsec used in the new DFT processor, L = 113, which gives a coefficient spacing of approximately 1 MHz.

The new DFT processor, as presently programmed, first computes 45 DFT coefficients spaced 1 MHz from 2 to 46 MHz. This is accomplished by evaluating Eq. (13) for k = 82 to 1,886 in steps of 41. After finding the locations of the tallest peak, k_p, and the second tallest peak, k_s, 9 more coefficients around the tallest peak are computed. These second-pass coefficients are obtained by evaluating Eq. (13) for $k = k_p - 32$ to $k_p + 32$ in steps of 8. This spans the frequency range of ± 0.78125 MHz around the first-pass peak, covering more than the ± 0.5 MHz uncertainty of the first pass, and so guaranteeing that the second-pass peak does not fall at the beginning or end of the second-pass coefficients. This, in turn, guarantees the existence of a coefficient to both the left and right of the second-pass peak as required for parabolic interpolation (Ref. 1).

Prior to outputting the peak information to the host computer, the DFT processor validates the spectrum and associated information. This is done by comparing the height of the second-tallest first-pass peak to the height of the second-pass peak. If the height of the second-tallest peak is greater than 50 percent of the height of the tallest (second-pass) peak, the spectrum and measurement are rejected as having noise peaks too near the same order-of-magnitude as the signal peak. Otherwise, the spectrum and peak information are validated and fed to the host CPU where the parabolic interpolation is performed to obtain the final refined estimate of signal frequency.

The final coefficient spacing of the new DFT processor using the algorithm above is 0.195 MHz, which is twice as good a resolution as that obtained with the prototype processor. It should be pointed out that the new processor is not limited to this resolution prior to the parabolic interpolation. In fact, from Eqs. (12) and (8) with M = 4,096, we find that the fundamental resolution is 24.4 kHz, which is eight times better than could be obtained with the prototype DFT processor.

3.0 SYSTEM HARDWARE

A block diagram of the production DFT LV system hardware is shown in Fig. 4. The transient digitizer was developed specifically to sample LV signals and quickly transfer signal samples

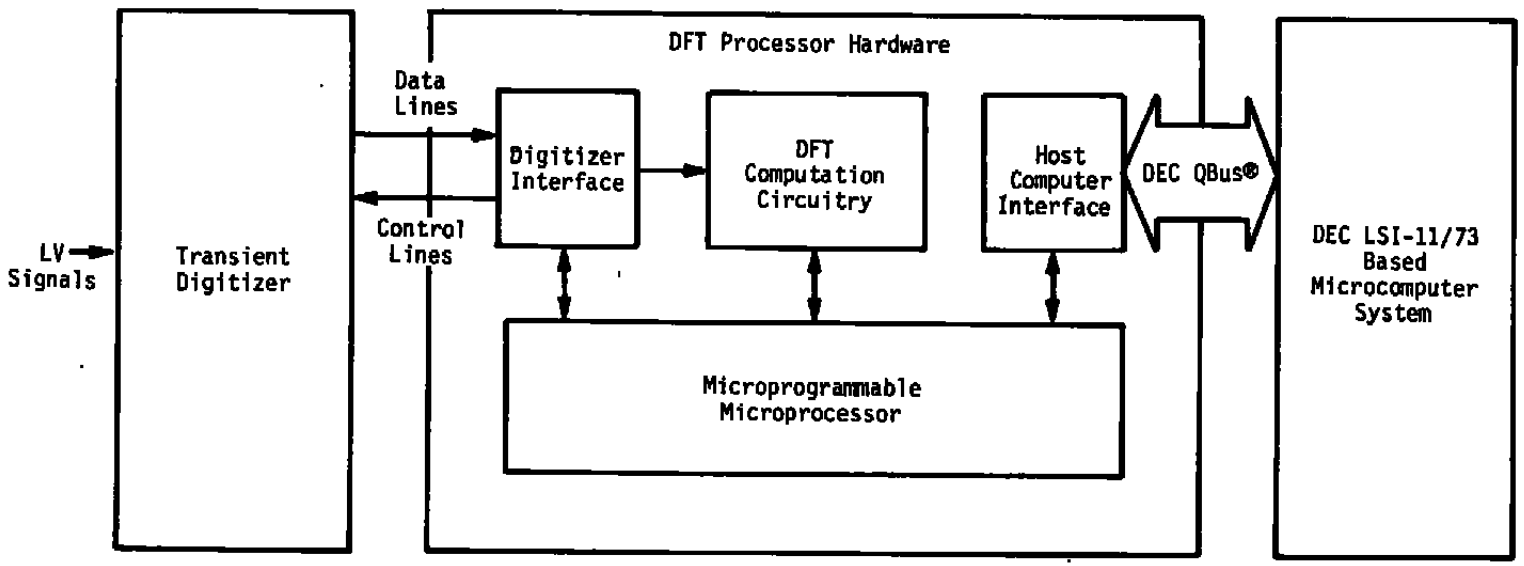

Figure 4. DFT processor system hardware.

to the DFT processor for spectral analysis. At the core of the DFT processor is a microprogrammable microprocessor controlling circuitry specially designed for fast, efficient, fixed-point-arithmetic calculation of the discrete Fourier transform. Spectral data derived from the analysis of LV signals are passed to the host computer where spectral peaks, corresponding to single-particle velocity measurements, are identified and stored. The DFT processor hardware resides on a single printed circuit board, which plugs directly into the backplane of the host computer system. Based on the Digital Equipment Corporation's LSI-11/73 CPU, the computer system communicates with the DFT processor board via the computer's internal QBus®.

3.1 WAVEFORM DIGITIZER

To avoid having the new DFT processor's data rate limited by the digitizer's low waveform transfer rate (as it was in the case of the prototype processor), a digitizer was designed in-house for use with the DFT processor. The transient digitizer was designed to meet the requirements listed as follows:

1. Input signals must be digitized to a 6-bit resolution at a sample rate of 100 MHz, corresponding to a 50-MHz limit on input signal frequency.

2. A sampled waveform will consist of 128 points spaced 10 nsec apart, defining a 1.27-μsec time-domain window.

3. The digitizer must be gated on by an external trigger input; this input is connected to some signal detection circuit so that only legitimate LV signals are digitized.

4. Provision must be made to store signal samples taken prior to the occurrence of the digitizer trigger pulse; by triggering the digitizer near when the expected peak signal level occurs and storing the 64 samples taken just prior to triggering as well as the 64 samples following the trigger pulse, the complete signal burst can be captured in the center of the time-domain window without having to set the trigger level down near the background noise.

5. The digitizer must be able to digitize three analog inputs simultaneously, allowing two- and three-component velocity measurements.

6. Waveform transfer time must be negligible compared to the DFT processor's waveform processing time, allowing the processor to achieve its maximum data rate.

The transient digitizer, shown in Fig. 5, is based on a commercially available 6-bit, flash analog-to-digital converter (A/D) integrated circuit, TRW's TDC1029. A 100-MHz oscillator drives the A/D's convert line to achieve a 10-nsec sample spacing, Δt. The A/D is allowed to run freely, continuously sampling its input. Signal samples are written into a first-in, first-out-type (FIFO) buffer memory where the most recent 128 samples are retained. When the external trigger input is activated, the posttrigger sample counter allows only N_{post} more samples to be written into the buffer before disabling the buffer's write line and activating the waveform ready line. Thus, $(128 - N_{post})$ samples are taken prior to the occurrence of the trigger pulse, and N_{post} samples are taken after the trigger. N_{post} can be varied from 2 to 128 samples to accommodate various triggering schemes.

The DFT processor independently accesses each channel of the digitizer via the A/D address bus. Once an address is set on the bus, the data strobe line is toggled to transfer signal samples from the channel addressed to the DFT processor. This interface allows 128-point waveforms to be moved to the processor at a rate in excess of 18,000 waveforms per second. The A/D reset line is common to all channels, enabling the processor to simultaneously reset all channels of the digitizer.

Presently, the A/D bus is daisy-chained as shown in Fig. 5. In this configuration, a single DFT processor board processes waveforms from all active digitizer channels. However, by splitting the A/D bus so there is a DFT processor board for each active digitizer channel, data rates for two- and three-component systems can be increased by factors of two and three, respectively.

3.2 PROCESSOR HARDWARE

Figure 6 is a detailed block diagram of the DFT processor. This diagram is divided into four functional areas described in the following sections.

3.2.1 Microprogrammed Microprocessor

The controlling element in the DFT processor is a microprogrammed microprocessor. This microprocessor consists of an Advanced Micro Devices (AMD) Am2910A microprogram sequencer, an AMD Am29116 16-bit arithmetic logic unit (ALU), and a 2,048-word by 64-bit microprogram memory composed of eight AMD Am27S45 programmable read-only memories (PROM's). The Am2910A sequencer controls not just the ALU, but all other components in the DFT processor as well. This sequencer was selected because of its ability to execute a microinstruction in only 0.1 μsec.

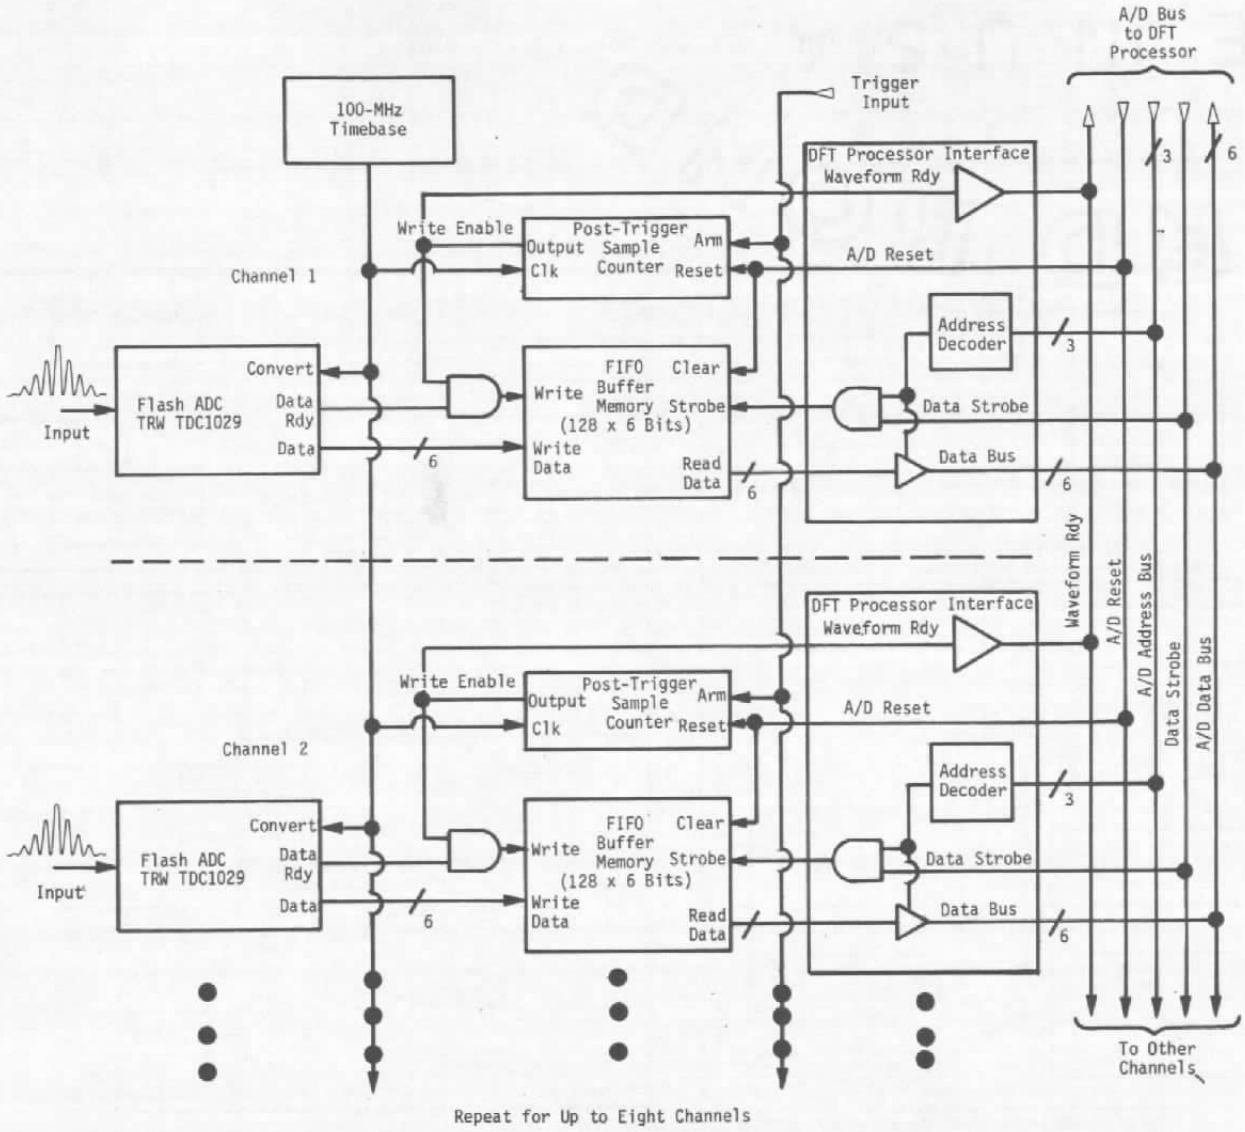

Figure 5. Transient digitizer block diagram.

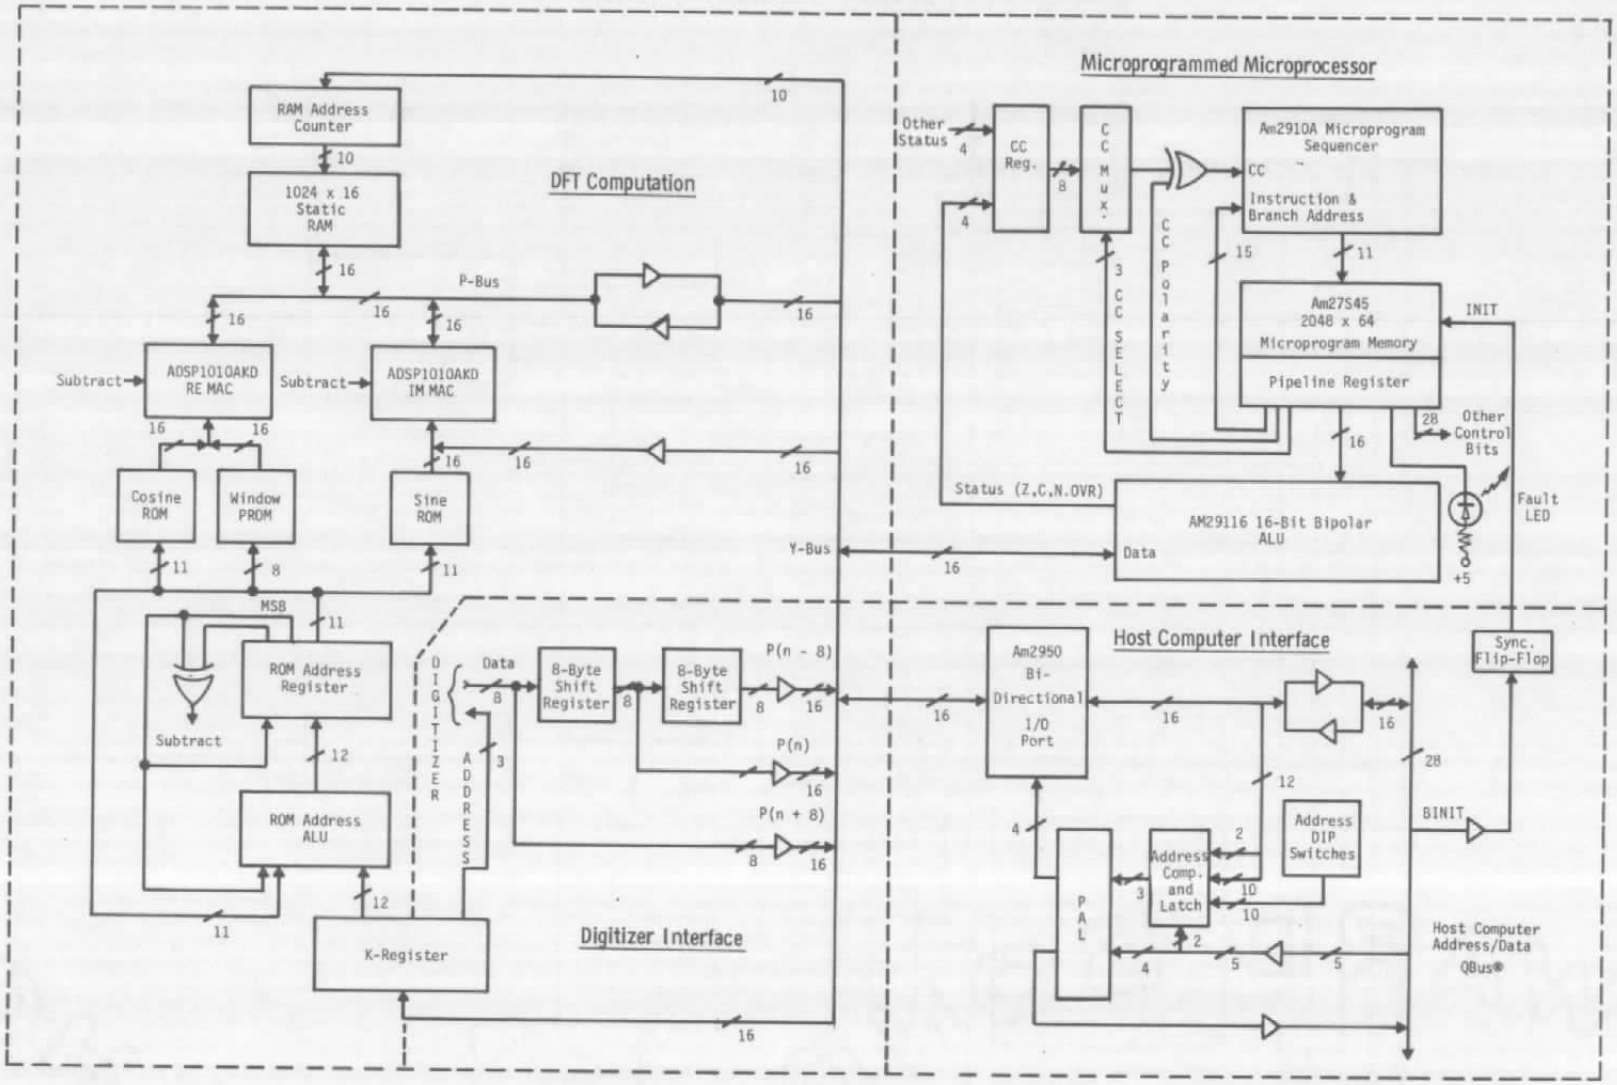

Figure 6. DFT processor block diagram.

The Am29116 ALU was chosen because it represents the latest in high-speed VLSI technology. This ALU occupies a single integrated circuit, as compared to four integrated circuits for equivalent functions in the prototype DFT processor (Ref. 1). The Am29116 is capable of executing one instruction every 0.133 μsec. To permit full utilization of both the sequencer and ALU, the ability to switch clock periods under microprogram control was included in the DFT processor. When the ALU is used during a microinstruction, the clock period is set to 0.133 μsec. Otherwise, the clock period can be set to 0.1 μsec.

3.2.2 Digitizer Interface

The digitizer interface is composed primarily of tristate buffers that permit digitizer data to be read onto the DFT processor's internal Y-bus. Two variable-length (each nominally configured as shown for an 8-byte length) shift registers were included to permit older values of the data to be retained for use by the trend-removal filter algorithm. The interface also provides three address bits, which specify the digitizer channel from which data are desired in multicomponent systems.

3.2.3 Host Computer Interface

The DFT processor was designed to work with any host computer in the Digital Equipment Corporation LSI-ll family. The present highest-performance member of this family, the LSI-11/73, is currently used with the DFT processor.

The DFT processor was designed to operate as an input/output (I/O) port for the host computer. As such, it communicates directly with the host computer address/data bus (the LSI-11 QBus). This is accomplished by assigning to the DFT processor four memory locations in the upper 4,096 words of computer memory (the LSI-ll I/O page). The particular locations assigned are determined by miniature switches on the DFT processor board. The four memory locations make up the input control/status register, the output control/status register, the input data buffer register, and the output data buffer register, which are standard for most bidirectional LSI-11 interfaces.

The host computer interface hardware consists primarily of two AMD Am2950 bidirectional I/O port integrated circuits, a programmable array logic (PAL) integrated circuit, and tristate buffers used to place data on the Y and Q buses. The single PAL integrated circuit contains all of the individual gates needed for the interface.

3.2.4 DFT Computation Circuitry

The heart of the DFT computation circuitry is two, 16-bit multiplier-accumulators (MAC's) used to calculate the products and sums-of-products needed for windowing and DFT-coefficient calculation. High-speed CMOS MAC's manufactured by Analog Devices (Model ADSP 1010 AKD) were used. Although these MAC's are pin-for-pin compatible with those used in the prototype DFT processor, the CMOS MAC's are more than twice as fast and consume several times less power. These MAC's permit a product to be formed and accumulated into a sum in the 0.1-μsec cycle time of the microprogram sequencer.

A high-speed 1,024-word by 16-bit static random access memory (RAM), composed of two Mostek MK4801AP-70 integrated circuits, is used to store digitizer data and DFT coefficients. The speed of this memory permits an address counter to sequentially access a new data sample every 0.1 μsec as required by the MAC's. The other data sources for the MAC's are read-only memories (ROM's) containing tabulated values of the sine and cosine functions, and a PROM, which contains values for the Hamming window function. Two AMD Am27S29 PROM's are used to hold the Hamming window values, and AMD Am29526/27 and Am29528/29 sine and cosine ROM's are used.

To access the sine and cosine ROM's in the manner necessary to compute the DFT at the top rate of the MAC's, a dedicated ALU is used to generate the ROM addresses. Since this ALU only needs to perform add, pass, and clear functions, relatively simple Am25LS2517 integrated circuits are used.

3.3 ALGORITHM IMPLEMENTATION

The DFT processor algorithm has been described in Section 2. The following discussion will consider how this algorithm is implemented using the hardware of Fig. 6.

3.3.1 Trend Removal and Windowing

The operations of trend removal and windowing performed as data are transferred from the digitizer to the DFT processor RAM. The DFT processor is set up for these operations by

1. loading the RAM address counter with the starting address in RAM where the trend removed and windowed data are to be placed;

2. selecting the window PROM as data source for the RE MAC; and

3. setting up the K-register, ROM address ALU, and ROM address register to sequentially address the window PROM (i.e., place 1 in the K-register, clear the ROM address register, and set the ALU to add).

The first 16 waveform samples are then read from the digitizer and summed in the Am29116 as required by Eq. (2). Simultaneously, these same samples are entered into the digitizer interface shift register. The outputs of the shift register then provide the values of p(n − 8), p(n), and p(n + 8) required for the AM29116 to compute Eq. (3). After each computation of Eq. (3), the trend-removed sample, q(n), is transferred from the Am29116 to the RE MAC along with the corresponding value of w(n) from the window PROM. The RE MAC then performs the product of Eq. (5) and the result is transferred to RAM. Each time Eq. (3) is evaluated, the new digitizer sample, p(n), is entered into the shift register in preparation for the next evaluation.

The trend-removal and windowing operations outlined above require a total of 4 microinstructions per sample to perform. Thus, data from the digitizer can be read, trend removed, windowed, and transferred to DFT processor RAM at a rate of one sample every 0.53 μsec.

3.3.2 DFT Computation

Computation of DFT coefficients using Eq. (13) is performed efficiently by recognizing two characteristics of the equation. First, it is noted that the Eq. (13) can be written

$$X(k) = u - jv \tag{16}$$

giving

$$|X(k)|^2 = u^2 + v^2 \tag{17}$$

where

$$\left.\begin{aligned} u &= \sum_{n=0}^{L-1} x(n) \cos\left(\frac{2\pi}{M} kn\right) \\ v &= \sum_{n=0}^{L-1} x(n) \sin\left(\frac{2\pi}{M} kn\right) \end{aligned}\right\} \qquad k = 0, 1, \ldots, \frac{M}{2} \tag{18}$$

Thus, the values of u and v can be simultaneously calculated using the RE and IM MAC's, respectively.

Second, it is noted that values of the cosine and sine functions are periodic with period 2π. Thus, only the values of these functions between 0 and 2π need to be tabulated in ROM. From Eq. (18) it is clear that M sine and cosine values must be tabulated where, from Eq. (12), the value of M determines the basic processor resolution. For the sine and cosine ROM's used in the DFT processor, M = 4,096 giving a basic resolution of 24.4 kHz.

To compute the DFT coefficient of Eq. (17) for a specified value of k (i.e. for a specified frequency), it is first necessary to calculate u and v from Eq. (18). The DFT processor is set up for this calculation by

1. loading the RAM address counter with the starting address in RAM where the trend-removed and windowed data are stored;

2. clearing the ROM address register;

3. setting the ROM address ALU to the add mode; and

4. placing the specified value of k in the K-register.

The L waveform samples are then successively applied to one input of the MAC's at a rate of one sample each microinstruction cycle. At the same time, the necessary values of the sine and cosine functions are accessed from ROM and applied to the other input of the MAC's. The RE MAC then forms the u sum-of-products while the IM MAC forms v.

From Eq. (18) it is noted that the required ROM addresses are kn for n = 0, 1, ..., L − 1. This gives addresses of 0, k, 2k, 3k, ..., (L − 1)k. The initial zero address was established in Step 2 above. The other addresses are formed, as required, by adding k to the ROM address register each microinstruction cycle using the ROM address ALU.

Although values of kn can exceed the 12-bit address space of the sine and cosine ROM's, this only happens for arguments that exceed 2π. Since the ROM address ALU has only 12 bits, higher-order bits are ignored causing addresses outside the ROM address space to "wraparound" into the ROM address space. The resulting addresses still give the correct sine and cosine values because of the 2π periodicity of these functions.

There is one other characteristic of the sine and cosine functions that is used to simplify the DFT processor hardware. Values of these functions for arguments between π and 2π are the negative of the values obtained with arguments between 0 and π. Thus, only 2,048 function values for the 0 to π range are actually stored in ROM. The most significant address bit does not then go to ROM but is used instead to determine the polarity of the values obtained from ROM. Whenever this polarity changes, the sign of the accumulated sum is changed by placing the MAC's in the subtract mode for one microinstruction cycle. Successive products then continue to be accumulated until the next polarity change. Although this approach can result in an incorrect sign for u and v, from Eq. (17) it is seen that the signs of u and v are irrelevant.

Once values for u and v have been obtained, they are squared and summed in the IM MAC. The resulting 32-bit DFT coefficient is then stored in RAM, and computation of the next DFT coefficient is initiated.

Most of the time involved in computing a DFT coefficient is spent in performing the sums-of-products of Eq. (18). Since one product can be formed and summed every microinstruction cycle, and since this operation does not involve the Am29116 ALU, the time required to form the sums-of-products for one DFT coefficient is 0.1 L μsec. The overhead of setting up to compute a coefficient and of calculating Eq. (17) is 1 μsec, giving a total calculation time of 1 + 0.1 L μsec per coefficient. For L = 113, this is 12.3 μsec.

3.3.3 Peak Location and Spectrum Validation

Once the DFT coefficients have been computed and stored in RAM, location of the peak and spectrum validation are straightforward. These operations are performed by the microprogrammed microprocessor. Validated spectra result in the peak location and the left-of-peak, peak, and right-of-peak DFT coefficients being output to the host computer. Spectra that can not be validated result in a reject code being output to the host computer.

3.4 WAVEFORM PROCESSING TIME

Most of the time required to process a waveform goes into trend removal, windowing, and DFT computation. Considering the present value of N = 128 (L = 113), trend removal and windowing require 60 μsec. Since the computation of 54 DFT coefficients is presently performed at the rate of one every 12.3 μsec, the DFT takes 664 μsec. Allowing another 76 μsec or so for miscellaneous operations such as peak location, spectrum validation, and communication with the host computer gives a total of 800 μsec to completely process a waveform,

corresponding to a processing rate of up to 1,250 waveforms per sec. This rate is achieved without the need for the frequency tracking scheme employed in the prototype DFT processor.

4.0 PROCESSOR DIAGNOSTIC AIDS

Several diagnostic aids were developed for use with the DFT processor. Processor diagnostic routines that exercise 75 percent of the processor's hardware have been included in the on-board microprocessor's microcode. These routines are executed at power-up and every time the host computer system resets its bus. Should the processor fail to successfully execute any of the diagnostic programs, the on-board microprocessor turns on the fault LED (Fig. 6) signifying that a processor malfunction has occurred and halts at the end of the failed routine. Thus, the DFT processor cannot be put into the data acquisition mode without first verifying that most of the signal processing hardware is operational. In addition to the automatic self-test routines, there are also three diagnostic routines that can be executed upon command from the host computer system. When used in conjunction with the diagnostic hardware described below, these routines allow the operational status of the DFT processor to be completely verified.

4.1 AUTOMATIC SELF-TEST ROUTINES

The automatic self-test routines primarily exercise the on-board microprocessor and the DFT computation circuitry. These routines occupy the first 184 words of microprogram memory and are executed in the order outlined in the following sections.

4.1.1 Microprocessor Diagnostics

First, the microprocessor diagnostic routines perform elementary operations in the microprocessor's arithmetic logic unit that should set or clear bits in the Am29116's status register. The microprogram sequencer (Am2910A) polls the state of the status bit under test to verify that it assumed the correct state for the operation that was performed. ALU status bits corresponding to negative result (N), zero result (Z), overflow (OVR), and carry (C) are tested by these routines.

Next, the Am29116's 32 internal RAM locations are checked by loading a negative one [or $(FFFF)_{16}$] into a given register, adding one to the contents of the register, and then checking for a zero result. Although these are not complete tests, they will generally be sufficient to identify a malfunctioning microprocessor.

4.1.2 DFT Computation Circuitry Diagnostics

After the microprocessor has successfully executed its diagnostic routines, each component of the DFT computation circuitry is tested. The 1,024-word by 16-bit RAM array used for storing digitized waveforms and DFT coefficients is tested first. Known bit patterns are stored in each location in the array, then the array is read to verify that it contains the correct data. Next the multiplier-accumulators that calculate the products and the sum-of-products needed for windowing and DFT-coefficient computation are exercised. Operations are performed that should yield a known result in the MAC's if they are operating properly. Finally, the read-only memories that contain the look-up tables for the sine, cosine, and window functions are tested. Three check sums are formed by totaling the contents of each ROM. The diagnostic program uses these calculated check sums to verify that the look-up tables are operational.

4.2 DIAGNOSTIC ROUTINES CALLABLE BY THE HOST COMPUTER

Presently the DFT processor can execute five different commands under control of the host computer. Of the five commands, two are intended primarily for data acquisition, and the other three are used for diagnostic purposes. The diagnostic commands provide a thorough checkout of the host computer interface, the digitizer interface, and the DFT processor as a whole.

4.2.1 Host Computer Interface Diagnostic

The host computer interface diagnostic command moves data from the DFT processor's input data register to its output data register. This allows the host computer to input a test data word into the DFT processor and then read it back to verify that the QBus interface is functioning properly. The other two diagnostic commands require the use of a circuit that emulates the waveform digitizer and supplies the DFT processor with a repeatable source of signal samples. The digitizer simulator uses an EPROM to store test waveform samples that are used in the checkout of the DFT processor.

4.2.2 Digitizer Interface Diagnostic

The digitizer interface diagnostic command allows unprocessed signal samples to be transferred from the waveform digitizer to the host computer. By transferring a test waveform from the digitizer simulator to the host computer, the digitizer interface can be tested with a precisely known waveform. When a calibrated analog input is connected to the waveform digitizer, this command also provides a means of testing the digitizer.

4.2.3 Integrated Diagnostic

The diagnostic aids discussed previously allow the functional areas that comprise the DFT processor to be individually tested. As the final step in processor checkout, an integrated diagnostic was developed that allows a processor to be tested as a whole. The integrated diagnostic allows the intermediate results, as well as the final results, obtained in processing a waveform to be checked against results obtained from a software-implemented version of the DFT processor.

The waveform digitizer simulator is used to provide a test waveform for the diagnostic. Once the test waveform has been processed, final results are stored in the on-board microprocessor's registers and results from intermediate steps in the processing algorithm reside in the DFT computation circuitry's RAM array. Upon command from the host computer, the contents of the microprocessor's registers and the RAM array are transmitted to the host system. These data are then compared to the same data obtained by running the computer simulation of the DFT processor. If the processor is functioning properly, it will generate results that are identical to those obtained from the simulation.

4.3 PROCESSOR DIAGNOSTIC BOARD

The processor diagnostic board was designed to aid in the repair of malfunctioning DFT processor boards. In designing the DFT processor board, provision was made to permit easy connection of the diagnostic board. The diagnostic board provides a means of monitoring, and optionally controlling, the DFT processor's on-board microprocessor. LED readouts on the diagnostic board display the current instruction in the pipeline register (Fig. 6) and the address of the next instruction to be executed. The DFT processor's master clock signal can be controlled via the diagnostic board. The clock can be allowed to freely run (normal mode of operation), single-stepped with a push-button switch, or halted at some predetermined breakpoint in the microprogram. The breakpoint is set with a group of miniature switches. Also included on the diagnostic board is an array of EPROM that may be configured to overrride the microprocessor's microprogram memory. Thus, special routines used in the repair of faulty processor boards can reside in EPROM on the diagnostic board, saving more of the processor's microprogram memory for the data acquisition code.

5.0 PROCESSOR EVALUATION

An experiment was devised that compares the DFT processor and the counter-type processor presently in use at the AEDC (Ref. 6) based on their ability to process noisy signals. Both processors were fed a signal consisting of a 15-MHz sine wave mixed with band-limited

white noise. The amount of noise added to the signal was varied to provide data points over a range of signal-to-noise ratios (SNR's). A block diagram of the experimental setup is shown in Fig. 7.

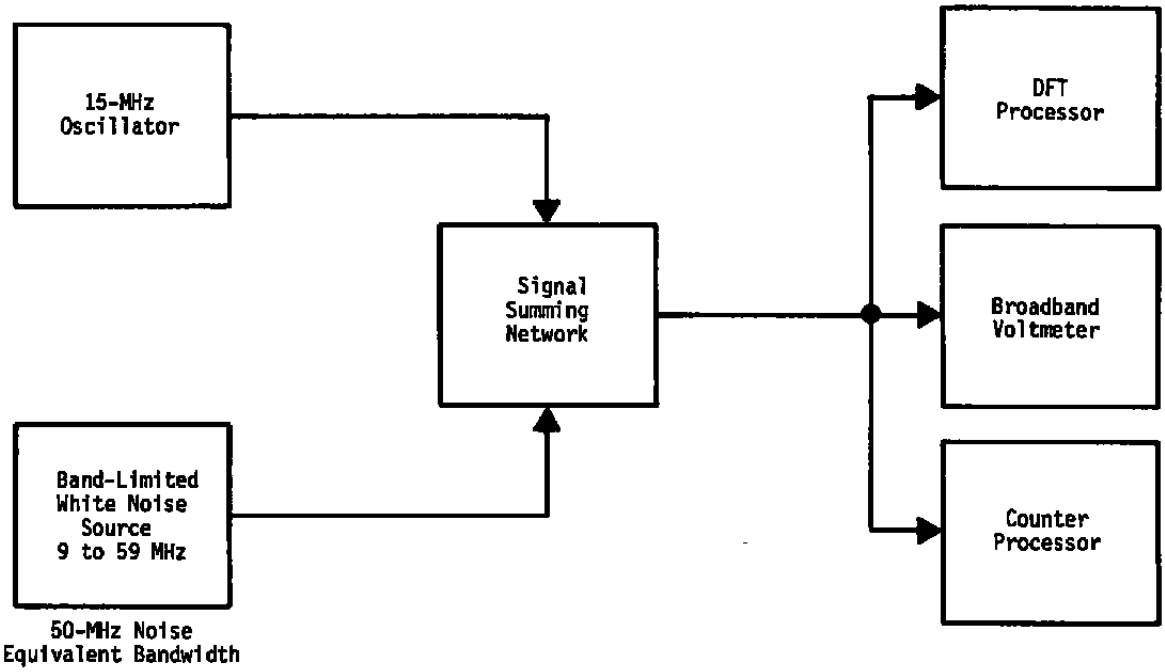

Figure 7. Signal-to-noise test setup.

Both signal processors reject erroneous data based on a predetermined measurement validation criteria. The counter processor required measurements based on four and five zero crossings to be within 3 percent of a measurement based on eight zero crossings (Ref. 6). The DFT processor required the amplitude of secondary spectral peaks to be less than half that of the primary peak. These are the measurement validation criteria typically used by LV signal processors at the AEDC.

For any given signal-to-noise ratio, each processor was allowed to take data until it accepted 1,000 readings. A data acceptance ratio was formed by dividing the number of validated readings (1,000) by the number of attempted readings. As shown in Fig. 8, the DFT processor was able to validate 100 percent of the signals it processed until the SNR fell below 1.1. On the other hand, the counter processor begins to reject signals at SNR's below 3.8. Figure 9 shows the percent error in the mean frequency measurement as a function of SNR. The DFT processor provides a good estimate of mean frequency until the SNR falls below 0.7, whereas the counter processor can give a similar estimate of mean frequency only when the SNR is above 1.6. The standard deviation for each 1,000 point sample is plotted in Fig. 10.

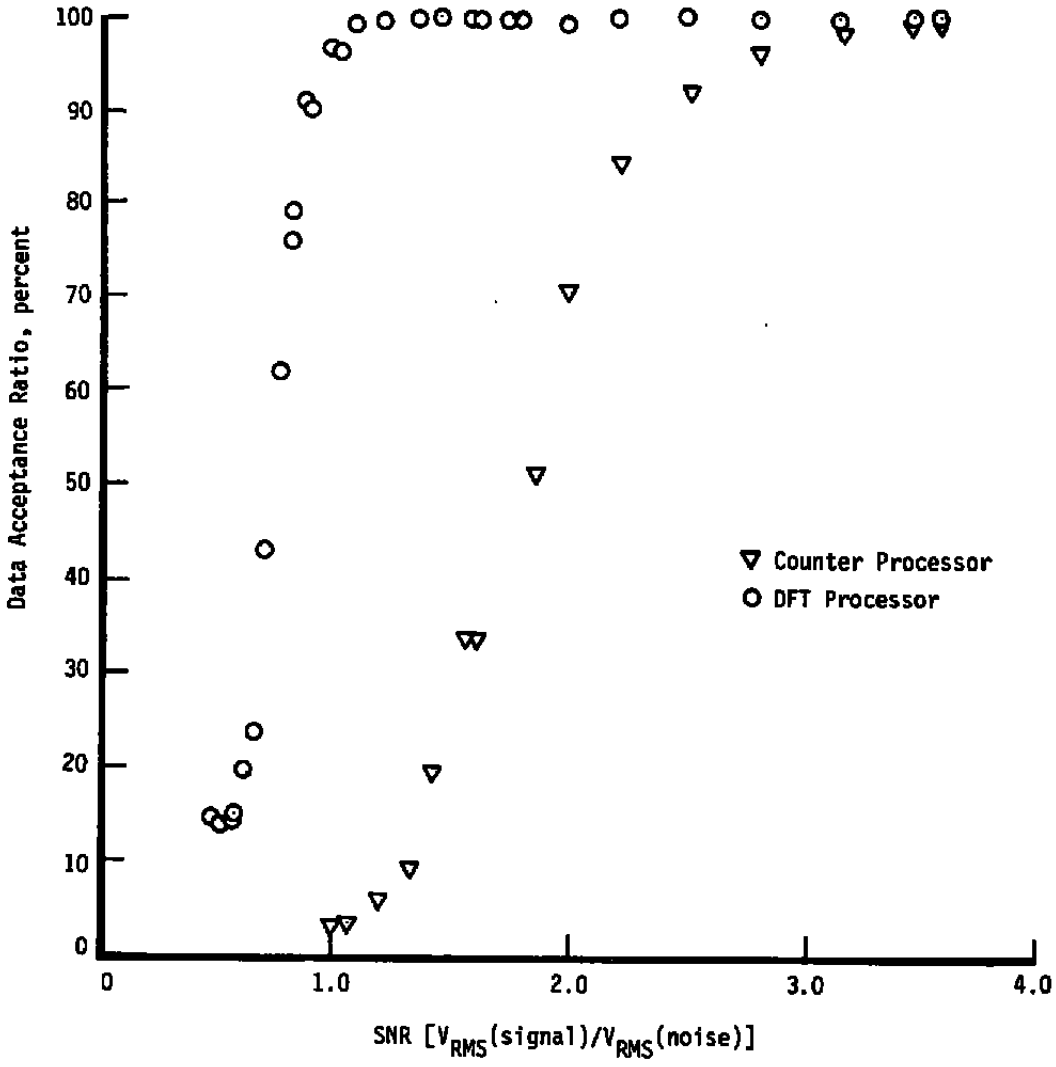

Figure 8. Data acceptance ratio plotted as a function of SNR.

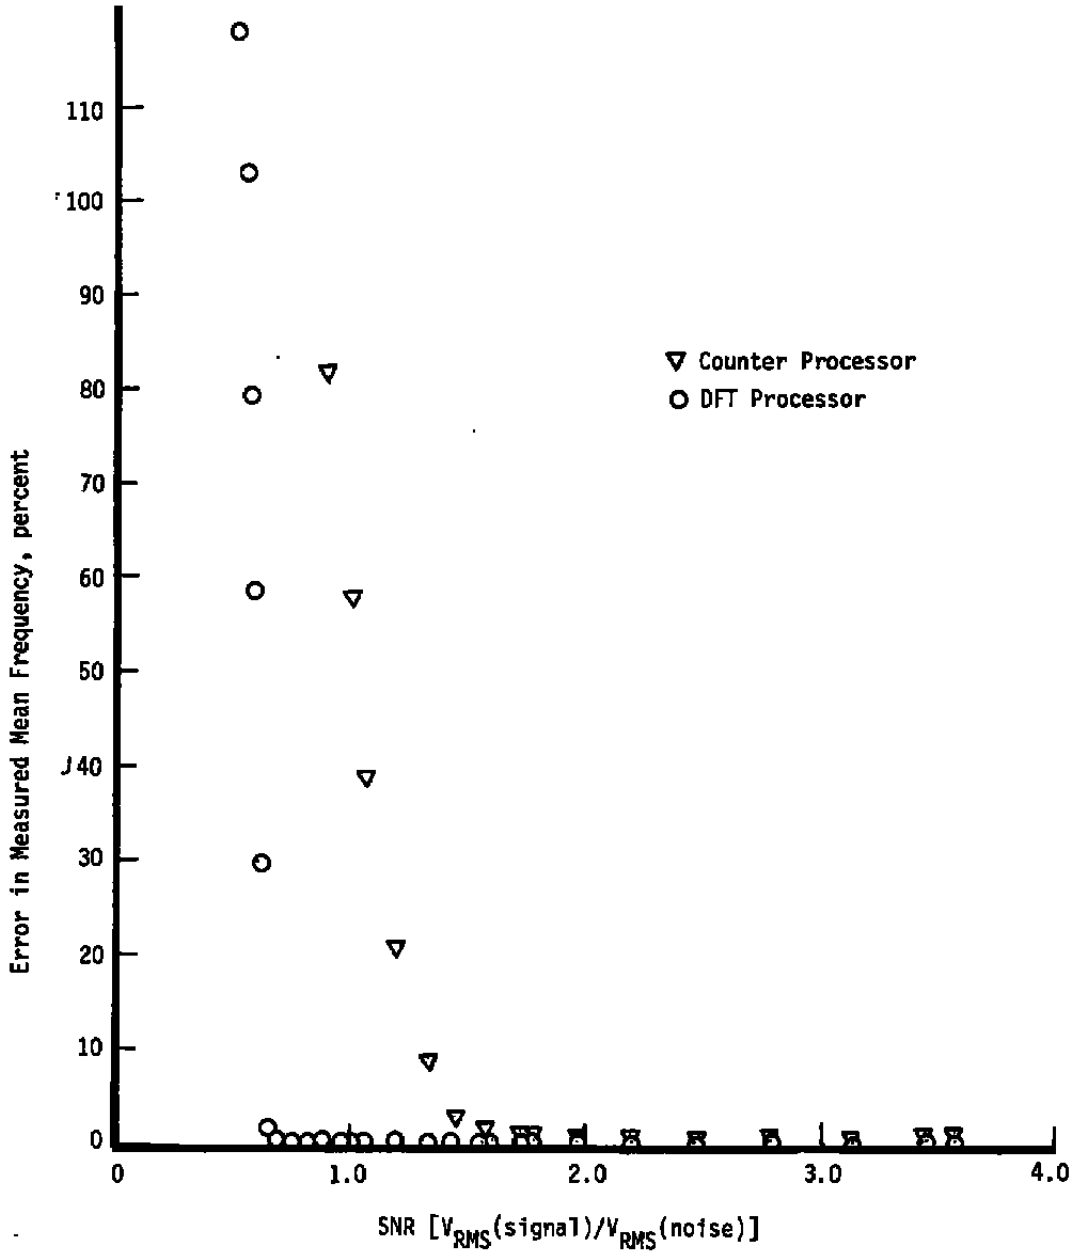

Figure 9. Percent error in measured mean plotted as a function of SNR.

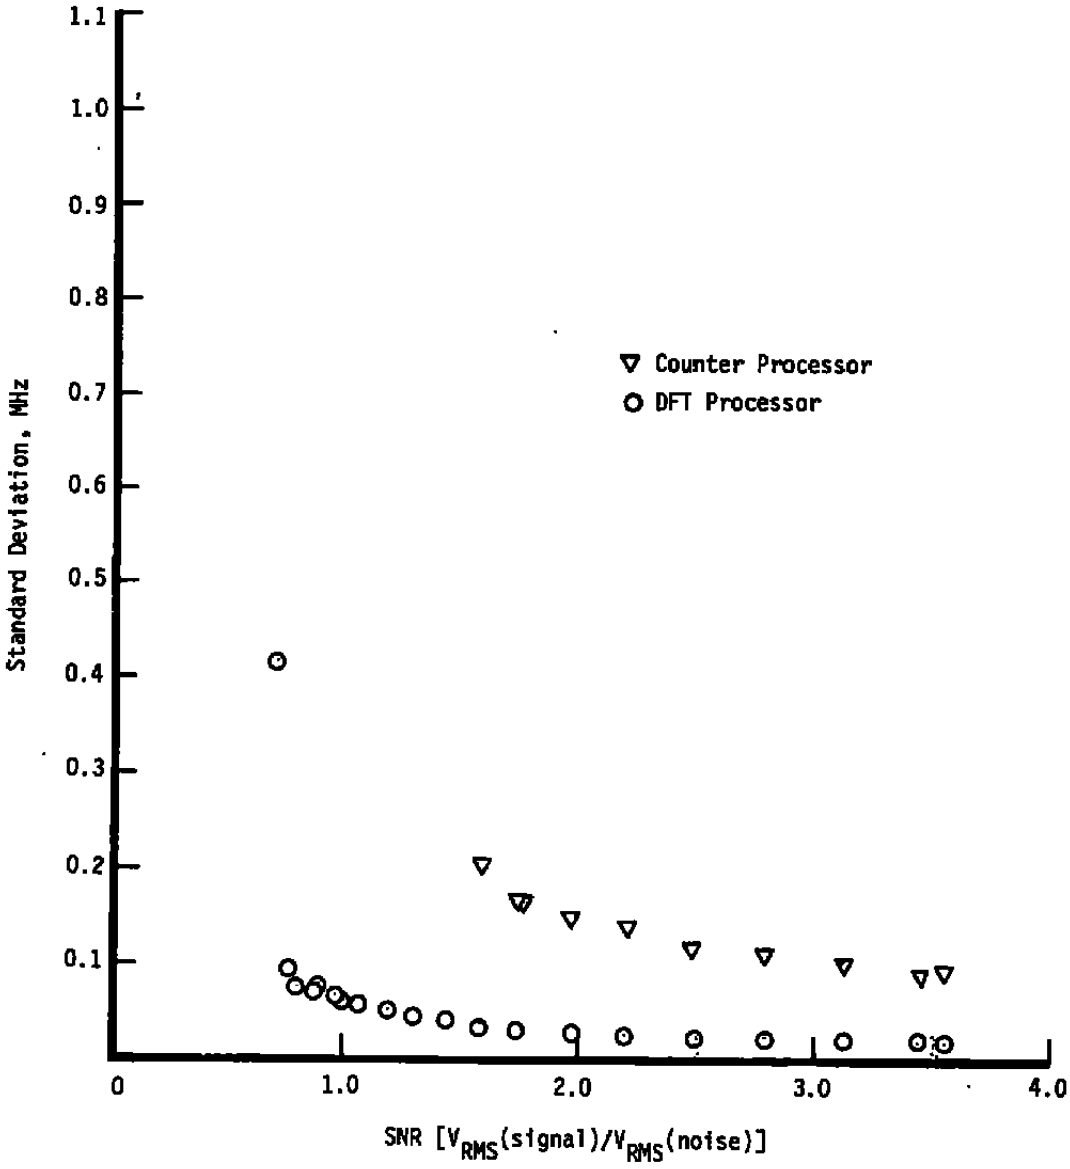

Figue 10. Standard deviation plotted as a function of SNR.

Before drawing conclusions based on Figs. 8, 9, and 10, the repeatability of these results was tested. Selected data points were repeated five times as shown in Table 1. For any given SNR the measured means and standard deviations vary only slightly from trial to trial. The data in Table 1 also agree well with Figs. 8, 9, and 10.

Table 1. Repeated Measurements of Selected Data Points Using 15.0-MHz Oscillator

Counter Processor

Signal-to-Noise Ratio				
Trial No.	0.667	1.019	1.533	2.093
1	33.045/5.770*	21.261/4.252	14.989/0.2159	14.984/0.1350
2	33.167/5.357	20.993/4.289	14.98/0.2021	14.989/0.1436
3	33.332/5.416	21.436/4.485	14.983/0.2203	14.984/0.1362
4	33.328/5.379	21.232/4.429	14.987/0.1945	14.987/0.1358
5	33.464/5.537	21.219/4.296	15.000/0.2505	14.992/0.1408

DFT Processor

Signal-to-Noise Ratio				
Trial No.	0.667	1.019	1.533	2.093
1	15.013/0.2933	15.00/0.059	14.999/0.0344	15.001/0.0252
2	15.0231/0.4306	14.997/0.0535	15.002/0.0331	14.999/0.0250
3	15.037/0.4522	14.999/0.0564	15.001/0.0325	15.9998/0.0250
4	15.018/0.3855	14.999/0.0542	15.002/0.0340	15.001/0.0244
5	15.003/0.1734	14.998/0.0528	14.999/0.0326	15.000/0.0247

* All table entries are mean/standard deviation.

6.0 SUMMARY, CONCLUSIONS, AND RECOMMENDATIONS

A second-generation DFT processor, based on Kalb's prototype (Ref. 1), has been designed, constructed, and evaluated. The availability of faster integrated circuits and improvements in the prototype processor's architecture and signal processing algorithm have made the production model DFT processor capable of a maximum data rate over six times greater than that of its predecessor. The increased data rate along with the inclusion of automatic self-test routines in the processor's microcode have transformed the prototype processor into a powerful, reliable tool for use in LV applications.

In the bench test presented in the previous section, the DFT processor consistently made measurements with less bias and precision error than the counter processor presently in use at the AEDC. The DFT processor was also superior in extracting measurements from poor-quality signals as illustrated in Fig. 8. It should be noted, however, that a counter processor that bases its measurements on more than eight cycles of the input signal would probably have less bias and precision error than the eight cycle counter. On the other hand, a counter that bases its measurements on more than eight cycles would probably have poorer results in the data acceptance ratio test.

Recent use of the new DFT processor during a test at Tunnel 1T demonstrated the ability of the processor to quickly extract high-quality measurements from poor-quality signals. The processor's performance in low signal-to-noise environments could be further improved by adding a spectrum-averaging mode to the processor's microcode. The prototype processor proved this mode of operation can extract measurements from extremely ill-defined signals that are individually unprocessable. In the spectrum-averaging mode, spectra of several signals are summed to form a composite spectrum before the spectral peak is located.

REFERENCES

1. Kalb, H. T. and Crosswy, F. L. "Discrete Fourier Transform Signal Processor for Laser Doppler Velocimetry." AEDC-TR-83-46 (AD-B078684L), December 1983.

2. Heltsley, F. L., Walker, B. J., and Nichols, R. H. "Transonic Nozzle-Afterbody Flow Field Measurements Using a Laser Doppler Velocimeter." AGARD CP-348, September 1983.

3. Heltsley, F. L. "Recent Experience in Seeding Transonic/Supersonic Flows at AEDC." NASA Conference Publication 2393, March 1985.

4. Oppenheim, A. V. and Schafer, R. W. *Digital Signal Processing*. Prentice-Hall, Inc., New Jersey, 1975.

5. Harris, F. J. "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform." *Proceedings of the IEEE*, Vol. 66, No. 1, January 1978, pp. 51-83.

6. Kalb, H. T., Brayton, D. B., and McClure, J. A. "Laser Velocimetry Data Processing." AEDC-TR-73-116 (AD-766418), September 1973.

NOMENCLATURE

A/D	Analog-to-digital converter
ALU	Arithmetic logic unit
CPU	Central processing unit
EPROM	Erasable-programmable read-only memory
FIFO	First-in, first-out
FIR	Finite impulse response
f_N	Nyquist frequency
f_s	Digitizer sample rate
I/O	Input/output
IM MAC	Imaginary component multiplier/accumulator
LED	Light emitting diode
LV	Laser velocimeter
MAC	Multiplier-accumulator
PAL	Programmable array logic

PROM	Programmable read-only memories
p(n)	Sampled photodetector output
p(t)	Photodetector output
q(n)	Digital filter output
RAM	Random access memory
RE MAC	Real component multiplier/accumulator
ROM	Read-only memory
VLSI	Very-large-scale-integration
x(n)	Windowed data
w(n)	Window function
ω	Digital radian frequency